DEALING WITH A LARGE INSTALLATION OF SRF CAVITIES: CHARACTERIZING LIMITATIONS AND EXPLOITING OPERATIONAL FLEXIBILITY*

C.E. REECE,† B.D. MADRE, L.R. DOOLITTLE,[1] and J.R. DELAYEN
Jefferson Lab, 12000 Jefferson Ave., Newport News, VA 23606, USA

Abstract

Linacs using a large number of SRF cavities can have an awkwardly large number of degrees of freedom for operational setup. The cost and robustness of operation as a function of operating gradient is a particular characteristic of each cavity system and the intended beamloading. A systematic characterization of these limitations has been developed which yields a valuable guide for development resource allocation. In addition, a software tool has been developed which enables the CEBAF machine operators to conveniently exploit the flexibility that results from the many degrees of freedom in response to changing programmatic needs. The two CEBAF SRF linacs each have about 160 independently-controlled SRF cavities. The software utility (LEM++) establishes the operationally optimum gradient in each cavity in response to the operator providing only three of the following four parameters: linac voltage, anticipated beam current, rf cryoheat load, and net rf trip rate. The utility is now fully operational at CEBAF. The methods employed and particular features useful for operations will be presented. The interactive process that has brought the software to its current form will also be discussed. The analysis scheme used to characterize the limitations of the ensemble of cavities will be presented as well.

1 BIG, STERILE PICTURE

Large-scale applications of SRF cavities very naturally push to obtain economical and efficient operation near their performance envelope, which is often defined by a complex interplay of multiple systems and "load" conditions. There is value to finding a convenient means of exploiting the many degrees of freedom.

The CEBAF recirculating electron linac was designed and built to provide 400 MV from each of two linacs, so that via five passes, one obtains 4 GeV electrons for nuclear physics research. The SRF systems in CEBAF have had no difficulty satisfying this requirement. However, as is commonly the case, the user community is eager to make use of the full range of capability of the machine, no matter what it is. To address this challenge requires a systematic understanding of various contributions to performance limitation and high-level "tuning knobs" which allow operations to approach the overall system envelope with the minimum degradation to accelerator reliability and availability.

2 BIG, REAL PICTURE

In the simplest of worlds, all of the installed SRF cavities meet specifications and are operated there; the R&D staff works on fundamental problems relevant to the next generation of applications.

In the real world, actual system performance varies, and the users want all that they can get, not just what was paid for. The most effective improvement strategies likely don't correspond to what seems most "interesting."

What to do?

- Characterize the actual system limitations
- Assess where lie the greatest opportunities for net improvement
- Allocate the limited R&D resources accordingly
- Develop tools that exploit all extant capacity
- Buy time and sneak (oh, we mean "leverage") resources necessary to make real progress

3 SOURCE DATA

Pushing the total system performance requires good characterization of the components. Cavity performance data collected during the 1991-1993 commissioning period served as the starting point for refinement of the limits. During the helium processing activities, each cavity was again pressed to its limits. All but the Q_0 data has been revisited. (Time did not permit new Q_0 measurements.) X-ray production was used as a surrogate indicator for Q degradation by field emission loading. Cavity quench fields were confirmed and the operational margin below quench reduced to 0.3 MV/m.

Acquiring good characterization of the window arcing behavior has required a significant amount of time. By running a few modules at a time at their limits for a few weeks, we accumulated adequate statistics with which to model the arcing rate. As a convenience, we defined as a standard limit for each cavity, that gradient which produces 3 trips per day. Although variation is observed, we model the variation of arc rate with gradient as increasing by a factor of 2 for each additional 0.35 MV/m near this reference gradient.

* Work supported by the U.S. DOE Contract # DE-AC05-84ER40150
† Email: reece@jlab.org

The historical Q_o values in the region of normal operation are used to optimize the linac voltage sum for the lowest 2 K cryogenic load.

4 LEM++

4.1 Optimization algorithm

Until recently, the cavity gradients were set up for operation by simply derating the setpoint of each cavity from its maximum by a uniform factor. This was adequate so long as there was generous capacity. The software utility that implemented this was Linac Energy Management (LEM). A new optimization solver, called LEM++, has been constructed that translates the 160 degrees of freedom (the gradient setpoints for each cavity), each with individual constraints, into just three.[2] Basically, we aim to set the operating gradient of each cavity as high as possible such that:

1. The linac voltage sums as needed.
2. The 2 K load/volt is no higher than necessary.
3. The cavity arcing rate/volt is no higher than necessary.
4. There remains adequate rf power for beamloading and regulation.

Phenomena considered include:

- Klystron output power limited to W Watts
- Cavity detuning $\delta = \tan\varphi$ (microphonics and static errors)
- Beam load from circulating current I
- Arc Trip Rate $R(V)$, modelled as $\exp(a + bV)$
- Cryogenic Load $C(V)$, modelled as V^2/R_d

Given constructive parameters I, $t_R \equiv \partial R/\partial V$, $t_C \equiv \partial C/\partial V$, compute

$$
\begin{aligned}
V(I) &= \frac{1}{1+\delta^2}\left[\sqrt{4WR_c(1+\delta^2) - I^2R_c^2\delta^2} - IR_c\right] \\
V(t_R) &= \frac{1}{b}\left[\ln\frac{t_R}{b} - a\right] \\
V(t_C) &= \tfrac{1}{2}t_C R_d \\
V_{\mathrm{SAFETY}} &= \text{constant} \\
V_{\mathrm{OPS}} &= \text{constant}
\end{aligned}
$$

The solution is the set of the lowest of these V that sums to the desired linac voltage. Allowance is provided in the voltage sum to reserve useful gradient range on four cavities in each linac that are used for beam energy stabilization and also to accommodate a non-linear sum due to less than perfect phasing of the cavities. This is the "fudge factor".

4.2 Operation

Each time LEM++ is used, the operator selects either the CEBAF North or South linac and provides three of the following four parameters: linac voltage, anticipated beam current, rf cryoheat load, and net rf trip rate. The solver finds the optimum value of the fourth parameter. An array on the user interface displays via color code the type of limit encountered for each cavity for the last solution. This provides useful qualitative feedback to the operator.

If the solution is appropriate for the upcoming program, the operator "applies" it, which initiates the following sequence:

1. Verification that adequate cryogenic capacity exists
2. Begin ramping the gradient setpoint of each cavity while pausing as needed to allow tuners to track and heat load allocation to shift between modules.
3. Loading linac quadrupole strengths to match the installed gradient profile.

The time required finding a solution and completing the application of the change ranges from 30 seconds to several minutes. The slowest part of the process is typically the tuner tracking rate, and this is normally significant only when large changes are made to cavities running above 9 MV/m and the Lorentz force tuning effect becomes significant.

The same process is used whether the operators are compensating for the failure of an individual klystron or making a large change in accelerator output energy.

4.3 User interface

Figure 1 shows the user interface for the North linac. In this example, the program calls for ~5.5 GeV. All 20 cryomodules are in use. For this high-energy setup, the operator solved for the minimum arc trip rate, having specified the linac voltage to be 550.0 MV, the expected total current envelope to be 300 µA and the rf cryoload to be 1275 watts.

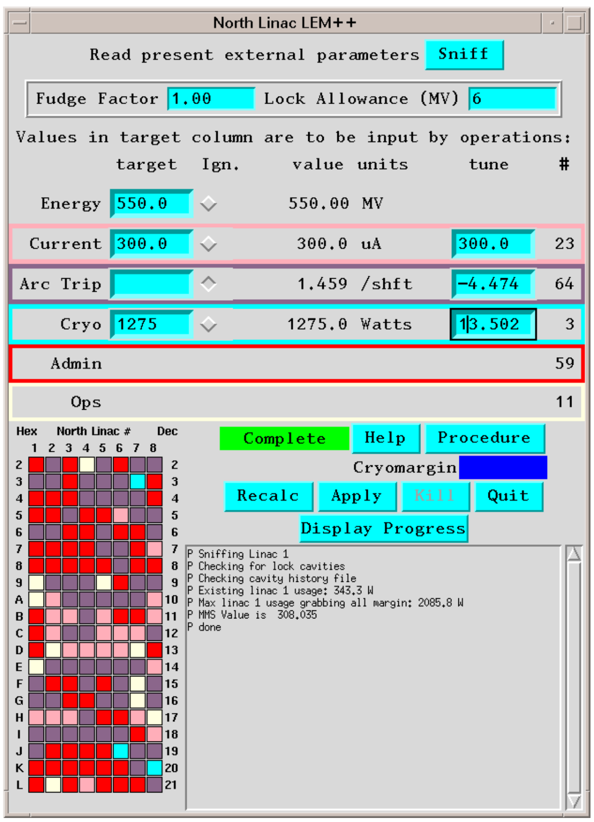

Figure 1. LEM++ user interface for the CEBAF North linac setup for 5.5 GeV.

With these constraints, about 1.5 arc trips per shift are unavoidable. Note that the 59 cavities with the “Admin” limit (e.g., limited by quench or field emission loading) are at their maximum, only three cavities are limited by their generated heat, “Cryo”, while 64 cavities contribute to the net arcing rate, and 23 are constrained by available rf power. LEM++ provides the operators with a means of applying their own temporary clamp due to short-term equipment problems. This is the case for the 11 cavities constrained by “Ops.”

4.4 Mapping options

Since the range of available solutions is *a priori* not at all apparent to the operators, the solver in LEM++ has also been used to map out the general solutions in the range of potential interest to the program. At different times, for example, cryogenic load may be no issue at all, while during other periods paying the price for additional refrigeration will translate into reduced arc trip rate and/or higher deliverable energy. See Figure 2.

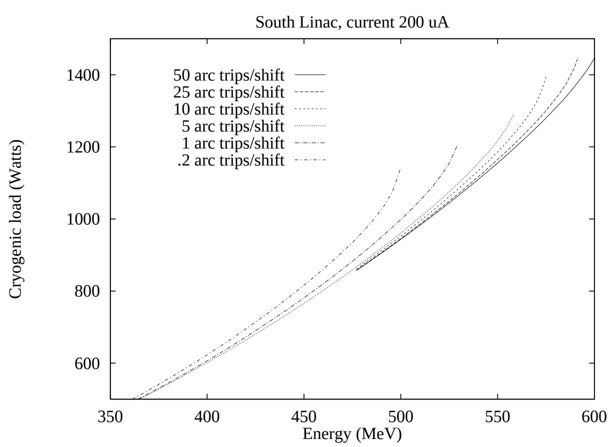

Figure 2. Range of available solutions corresponding to 5-pass energy of 3.5 to 6.0 GeV calculated by LEM++ for the CEBAF South linac with 200 µA total beam current.

4.5 Mechanics of the software

The software consists of four programs (modules) that can be tested independently. There is (1) the User Interface, (2) the Server, which calculates the cavity gradients, (3) linac_setup, which actually sets the gradients and quadrupole magnets, and (4) linac_sniff, which reads the status of the running machine. These programs pass information via ASCII lines which allows easy debugging and scripting. The server starts linac_sniff and linac_setup as batch jobs when requested through the user interface. Most of the communication is done through temporary files, but some "live" (STDOUT) messages come through as diagnostics and progress reports to the user.

The user interface and the server communicate bi-directionally via Unix pipes, but code exists to allow message passing with network sockets (not currently in use).

Messages have a consistent format:

<destination code> <function> <data>

with fields separated by white space. The function is case sensitive, and the data types and formats depend on the function. The destination code is a single letter. To illustrate:

C gset R221 4.030

This line is directed at the linac_setup module (its messages start with C), and tells it to set the process variable R221GSET to 4.030. Progress is indicated to the user through a window on the interface screen, shown above. The user can select the amount of information to be displayed in that window.

5 PERFORMANCE ANALYSIS MODEL

In addition to operational performance optimization, it is also useful to reduce the complexity of the multi-system performance parameters as one develops prioritized improvement strategies. Each cavity is constrained by one of several types of limits. (See Figure 3.)

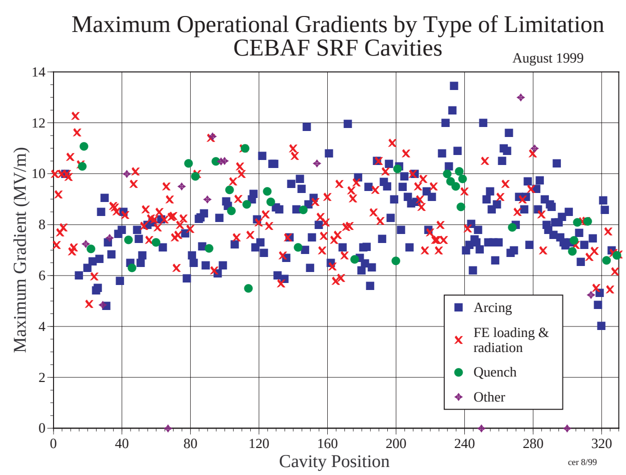

Figure 3. Maximum useful gradient and type of limitation for CEBAF cavities.

Distinguishing the importance of each type of limit may be difficult. We have developed an analysis model for performance limits of SRF cavities that allows us to determine the distribution of gradient that could be achieved for each type of limitation independently of the others.[3]

Three inputs are required to obtain the independent distribution functions, F_i:

1. The probability distribution of the actual gradient limits independent of limit type.
2. The cumulative rate of occurrence with gradient for each type of limit.
3. The assumption that the constraints are effectively independent of each other.

Figure 4 shows the calculated F_i in Summer 1998. For the purpose of guiding decisions, it is adequate to assume that the corresponding probability density functions, f_i, have an approximately log-normal form. The smoothed functions f_i (G) are shown in Figure 5, along with the overall density function f_i G).

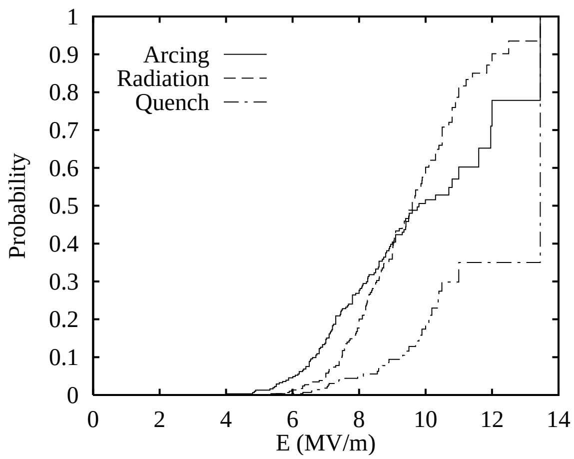

Figure 4. Calculated distributions of F_i vs. accelerating gradient.

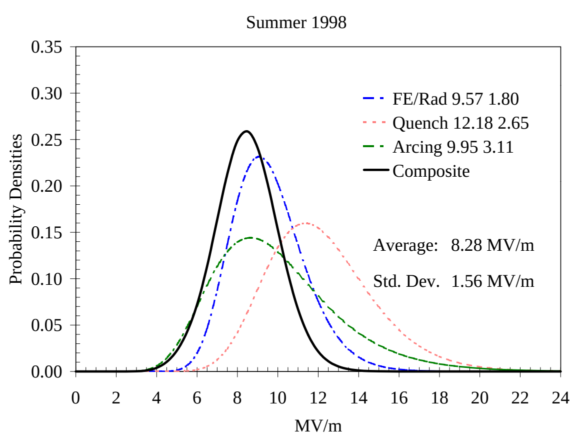

Figure 5. Probability densities for 1998 limitations.

This analysis is particularly helpful in identifying the type of improvements that would have the greatest impact on performance. For example, Figure 5, which represents the status of the CEBAF cavities in mid-1998, shows that their overall performance (labeled "composite" on the graph) could be most easily improved by a reduction of arcing (pushing the "Arcing" curve to higher gradients).

6 ACKNOWLEDGEMENTS

Several members of the CEBAF operations staff have provided valuable input to LEM++ so that it is conveniently useful to non-RF-experts. Great feedback has come from Jacque Ludwig, Anthony Cuffe, Michele Joyce, and Harry Fanning.

7 REFERENCES

[1] Present address: LBNL

[2] J.R. Delayen, L.R. Doolittle, and C.E. Reece, "Operational Optimization of Large-scale SRF Accelerators," Proc. PAC99, p. 940.

[3] J.R. Delayen, L.R. Doolittle, and C.E. Reece, "Analysis of Performance Limitations for Superconducting Cavities," Proceedings LINAC98, Chicago, p. 989.